



US009282032B2

(12) **United States Patent**
Judge et al.

(10) **Patent No.:** **US 9,282,032 B2**
(45) **Date of Patent:** ***Mar. 8, 2016**

- (54) **DISTRIBUTED ROUTING ARCHITECTURE**
- (71) Applicant: **Amazon Technologies, Inc.**, Seattle, WA (US)
- (72) Inventors: **Alan M. Judge**, Dublin (IE); **David J. McGaugh**, Seattle, WA (US); **James R. Hamilton**, Seattle, WA (US); **Justin O. Pietsch**, Bothell, WA (US); **David J. O'Meara**, Dublin (IE)
- (73) Assignee: **Amazon Technologies, Inc.**, Seattle, WA (US)
- (*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.
- This patent is subject to a terminal disclaimer.

- (21) Appl. No.: **14/632,962**
(22) Filed: **Feb. 26, 2015**

- (65) **Prior Publication Data**
US 2015/0172178 A1 Jun. 18, 2015

- Related U.S. Application Data**
- (63) Continuation of application No. 13/620,363, filed on Sep. 14, 2012, now Pat. No. 8,971,328, which is a continuation of application No. 12/641,260, filed on Dec. 17, 2009, now Pat. No. 8,331,371.

- (51) **Int. Cl.**
H04L 12/28 (2006.01)
H04L 12/58 (2006.01)
(Continued)

- (52) **U.S. Cl.**
CPC **H04L 45/44** (2013.01); **H04L 45/04** (2013.01); **H04L 45/58** (2013.01); **H04L 45/741** (2013.01); **H04L 45/745** (2013.01)

- (58) **Field of Classification Search**
CPC H04L 45/04; H04L 45/58; H04L 45/583; H04L 45/74; H04L 45/741; H04L 45/745
USPC 370/392, 400, 408
See application file for complete search history.

- (56) **References Cited**
U.S. PATENT DOCUMENTS

5,649,185 A 7/1997 Antognini et al.
6,182,125 B1 1/2001 Borella et al.
(Continued)

FOREIGN PATENT DOCUMENTS

CN 1511399 A 7/2004
CN 101473598 A 7/2009
JP 3748216 B2 2/2006

OTHER PUBLICATIONS

Al-Fares, M. et al., A Scalable, Commodity Data Center Network Architecture, SIGCOMM '08 Proceedings, Aug. 17, 2008, pp. 63-74, 66-68, 70-71, Seattle, WA.

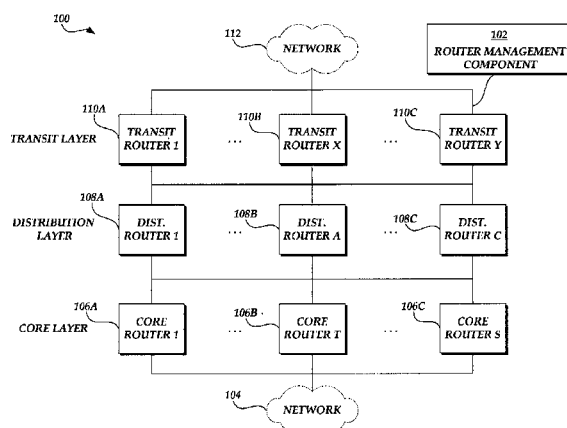
(Continued)

Primary Examiner — Dang Ton
Assistant Examiner — Ryan Kavleski
(74) *Attorney, Agent, or Firm* — Knobbe, Martens, Olson & Bear, LLP

- (57) **ABSTRACT**

A routing management component is provided for distributing routing information among a hierarchical distributed routing architecture. The routing management component can function to associate levels of the routing architecture with subsets of a network address format. The routing management component can further assign routers of the routing architecture to portions of network addresses defined at least in part by the network address format. For example, a router may be assigned to route packets addressed to a network address with a first octet between a range of values. The router management component may further distribute, to the routers of the hierarchical distributed routing architecture, sections of routing information associated with their assigned portions of network addresses. Because routing information can be distributed between various routers, the memory requirements of individual routers can be reduced comparatively to systems in which a single router maintains an entire set of forwarding information.

20 Claims, 7 Drawing Sheets



- (51) **Int. Cl.**
H04L 12/721 (2013.01)
H04L 12/715 (2013.01)
H04L 12/775 (2013.01)
H04L 12/56 (2006.01)
H04L 12/741 (2013.01)

(56) **References Cited**

U.S. PATENT DOCUMENTS

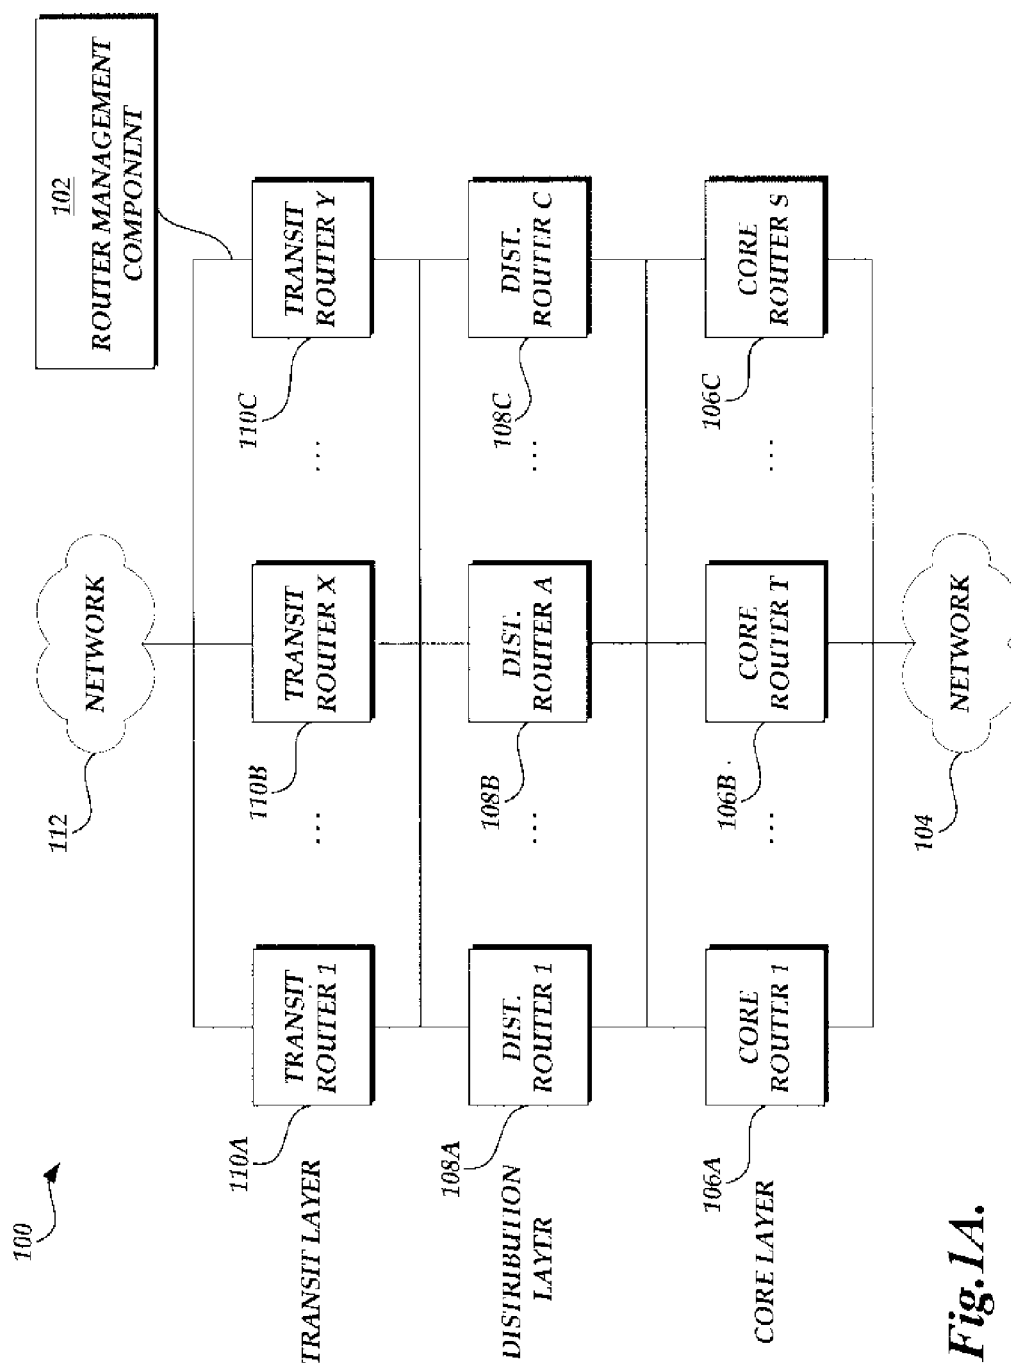
- 6,192,051 B1 * 2/2001 Lipman et al. 370/389
6,243,761 B1 6/2001 Mogul et al.
6,377,257 B1 4/2002 Borrel et al.
6,529,910 B1 3/2003 Fleskes
6,560,610 B1 * 5/2003 Eatherton et al.
6,633,324 B2 10/2003 Stephens, Jr.
6,697,805 B1 2/2004 Choquier et al.
7,009,943 B2 3/2006 O'Neil
7,023,465 B2 4/2006 Stephens, Jr.
7,065,496 B2 6/2006 Subbloie et al.
7,096,193 B1 8/2006 Beaudoin et al.
7,114,160 B2 9/2006 Suryanarayana et al.
7,120,874 B2 10/2006 Shah et al.
7,146,560 B2 12/2006 Dang et al.
7,185,084 B2 2/2007 Sirivara et al.
7,320,131 B1 1/2008 O'Toole, Jr.
7,502,836 B1 3/2009 Menditto et al.
7,523,181 B2 4/2009 Swildens et al.
7,555,542 B1 6/2009 Ayers et al.
7,623,460 B2 11/2009 Miyazaki
7,624,169 B2 11/2009 Lisiecki et al.
7,650,376 B1 1/2010 Blumenau
7,653,725 B2 1/2010 Yahiro et al.
7,676,570 B2 3/2010 Levy et al.
7,707,071 B2 4/2010 Rigole
7,707,173 B2 4/2010 Nanavati et al.
7,739,400 B2 6/2010 Lindbo et al.
7,773,596 B1 * 8/2010 Marques 370/392
7,787,380 B1 * 8/2010 Aggarwal et al. 370/236
7,865,594 B1 1/2011 Baumbach et al.
7,873,065 B1 1/2011 Mukerji et al.
7,904,875 B2 3/2011 Hegyi
7,930,393 B1 4/2011 Baumbach et al.
7,933,988 B2 4/2011 Nasuto et al.
7,937,456 B2 5/2011 McGrath
7,961,736 B2 6/2011 Ayyagari
8,051,166 B1 11/2011 Baumbach et al.
8,069,231 B2 11/2011 Schran et al.
8,117,306 B1 2/2012 Baumbach et al.
8,286,176 B1 10/2012 Baumbach et al.
8,296,429 B2 10/2012 Baumbach et al.
8,331,370 B2 12/2012 Hamilton et al.
8,902,897 B2 12/2014 Hamilton et al.
8,971,328 B2 3/2015 Judge et al.
2002/0135611 A1 9/2002 Deosaran et al.
2002/0138443 A1 9/2002 Schran et al.
2002/0150094 A1 * 10/2002 Cheng et al. 370/389
2002/0194382 A1 12/2002 Kausik et al.
2003/0149581 A1 8/2003 Chaudhri et al.
2003/0182413 A1 9/2003 Allen et al.
2004/0049579 A1 3/2004 Ims et al.
2004/0083307 A1 4/2004 Uysal
2004/0194085 A1 9/2004 Beaubien et al.
2005/0038967 A1 2/2005 Umbehocker et al.
2005/0086645 A1 4/2005 Diao et al.
2005/0102683 A1 5/2005 Branson et al.
2005/0157712 A1 * 7/2005 Rangarajan et al. 370/388
2005/0171959 A1 * 8/2005 Deforche et al. 707/100
2005/0259672 A1 * 11/2005 Eduri 370/412
2006/0209701 A1 9/2006 Zhang et al.
2006/0218304 A1 9/2006 Mukherjee et al.
2007/0021998 A1 1/2007 Laithwaite et al.

- 2007/0086429 A1 4/2007 Lawrence et al.
2007/0198982 A1 8/2007 Bolan et al.
2007/0219795 A1 9/2007 Park et al.
2007/0245010 A1 10/2007 Arn et al.
2008/0065724 A1 3/2008 Seed et al.
2008/0086559 A1 4/2008 Davis et al.
2008/0172488 A1 7/2008 Jawahar et al.
2008/0183721 A1 7/2008 Bhogal et al.
2008/0215755 A1 9/2008 Farber et al.
2009/0187575 A1 7/2009 DaCosta
2009/0248893 A1 10/2009 Richardson et al.
2009/0327517 A1 12/2009 Sivasubramanian et al.
2010/0128638 A1 * 5/2010 Navas et al. 370/254
2010/0150155 A1 6/2010 Napierala
2014/0129707 A1 5/2014 Baumbach et al.
2014/0143320 A1 5/2014 Sivasubramanian et al.
2014/0257891 A1 9/2014 Richardson et al.

OTHER PUBLICATIONS

- Bennami, M., et al., Resource Allocation for Autonomic Data Centers Using Analytic Performance Models, 2005, IEEE, 12 pages.
Chang, F., et al., Automatic Configuration and Run-time Adaptation of Distributed Applications, 2000, IEEE, 10 pages.
First Office Action in Chinese Application No. 201080057229.9 dated May 14, 2014.
Greenberg, A. et al., Networking the Cloud, 29th IEEE International Conference on Distributed Computing Systems (ICDCS 2009), Jun. 22, 2009-Jun. 26, 2009 [online] retrieved from the Internet on Mar. 10, 2011: http://www.cse.ohio-state.edu/icdcs2009/Keynote_files/greenberg-keynote.pdf, pp. 1-45.
Greenberg, A. et al., Towards a Next Generation Data Center Architecture: Scalability and Commoditization, SIGCOMM '08: Proceedings of the 2008 SIGCOMM Conference and Co-located Workshops NSDR '08, WOSN '08, MOBIARCH '08, NETECON '08, & Presto '08, Seattle, WA, Aug. 17-28, 2008, ACM, Aug. 17, 2008, pp. 57-62, New York, NY.
Greenberg, A. et al., VL2: A scalable and flexible data center network, SIGCOMM '09, Proceedings of the ACM SIGCOMM 2009 Conference on Data Communication, Aug. 17, 2009, vol. 39, Issue 4, pp. 51-62.
International Preliminary Report on Patentability and Written Opinion in PCT/US2010/060567 mailed on Jun. 19, 2012.
International Preliminary Report on Patentability and Written Opinion in PCT/US2010/060569 mailed Jun. 19, 2012.
International Preliminary Report on Patentability and Written Opinion in PCT/US2010/060573 mailed Jun. 19, 2012.
International Search Report and Written Opinion in PCT/US2010/060567 mailed on Mar. 28, 2012.
Kounev, S., et al., Autonomic QoS-Aware Resource Management in Grid Computing Using Online Performance Models, 2007, ICST, Valuetools, 2007, 10 pages.
Mysore, R.N. et al., Portland: a scalable fault-tolerant layer 2 data center network fabric, SIGCOMM '09, Proceedings of the ACM SIGCOMM 2009 Conference on Data Communication, Aug. 17, 2009, pp. 39-50.
Nilsson et al., IP-Address Lookup Using LC-Tries, IEEE Journal on Selected Areas of Communication, Jun. 1999, vol. 17, Issue 6, pp. 1083-1092.
Office Action in Canadian Application No. 2784699 dated Apr. 28, 2014.
Office Action in Canadian Application No. 2784706 dated May 22, 2014.
Office Action in Chinese Application No. 201080057155.9 dated Jul. 24, 2014.
Office Action in Chinese Application No. 201080057225.0 dated Jul. 2, 2014.
Second Office Action in Chinese Application No. 201080057229.9 dated Mar. 18, 2015.

* cited by examiner






Fig. 1B.

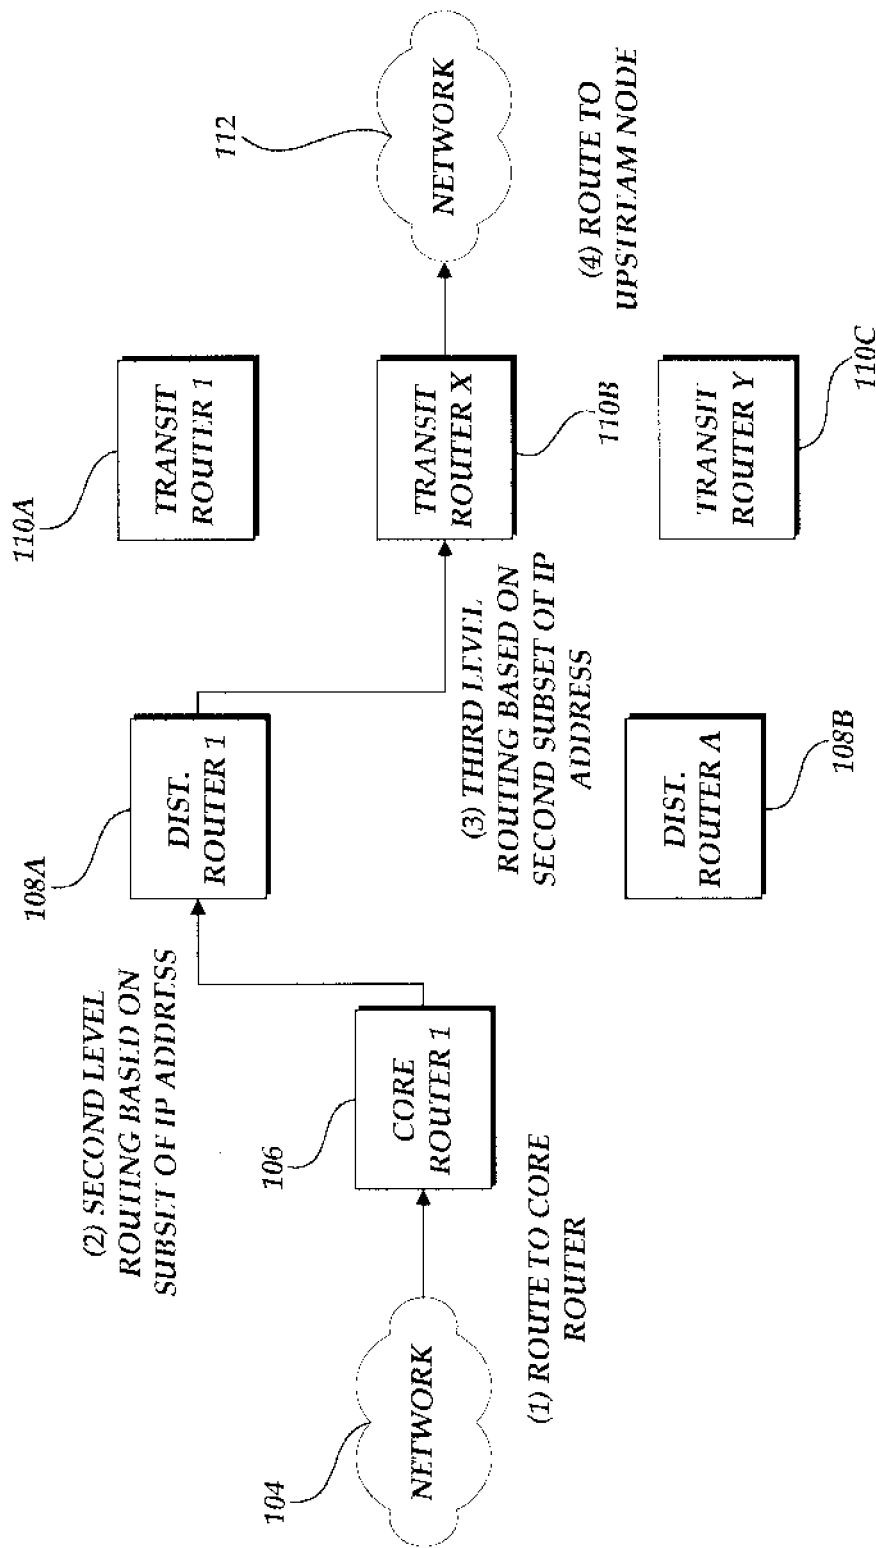


Fig. 2A.




Fig. 2B.



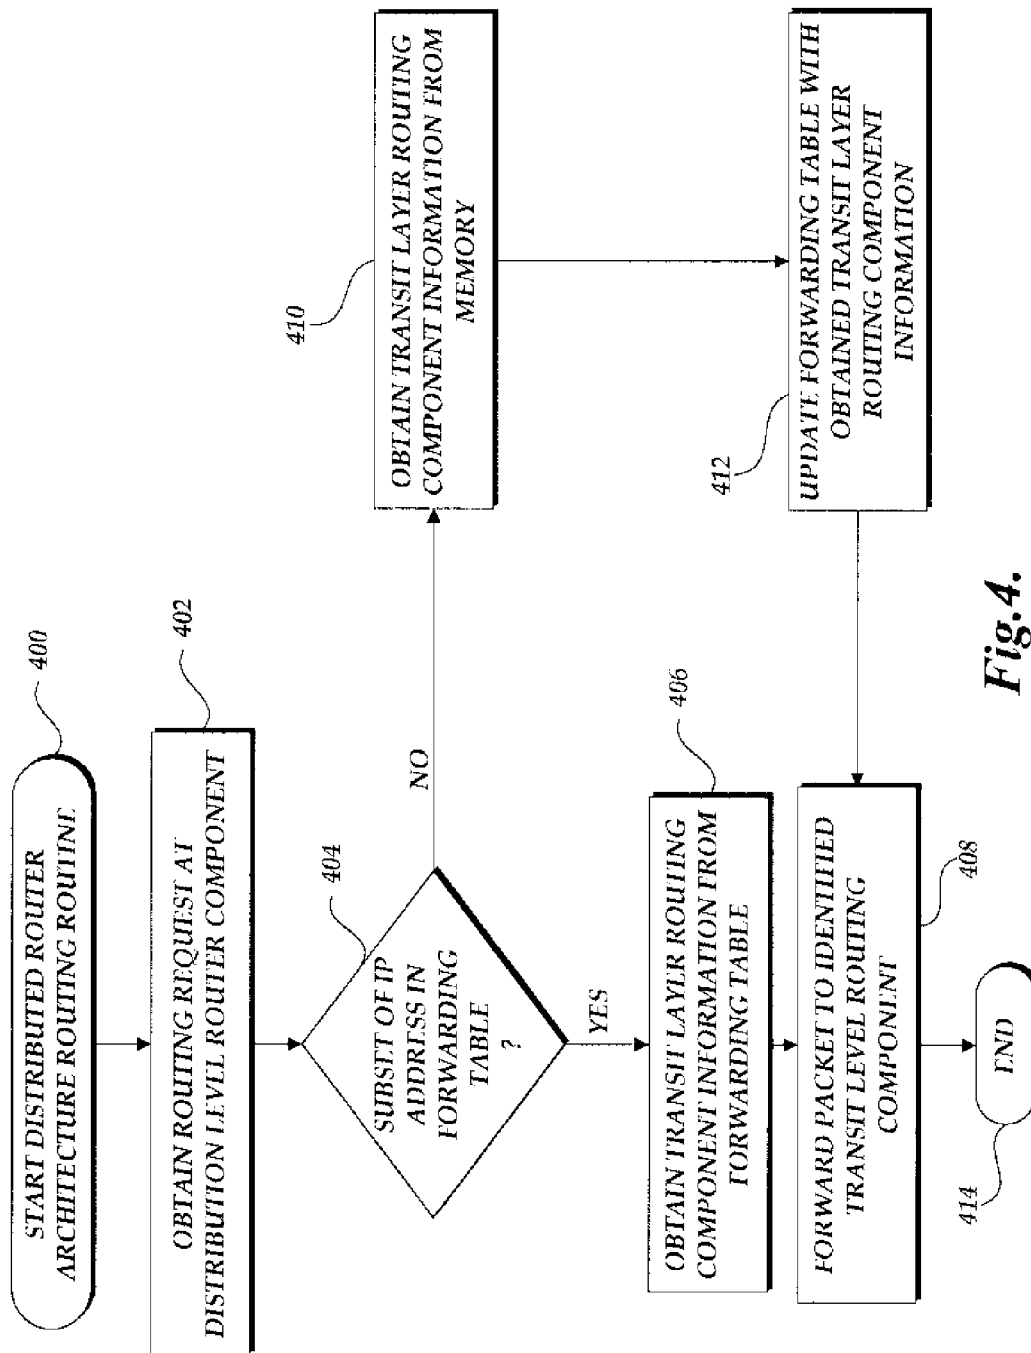


Fig. 2C.

**Fig.3.**

**Fig. 4.**

1

DISTRIBUTED ROUTING ARCHITECTURE**CROSS-REFERENCE TO RELATED APPLICATIONS**

This application is a continuation of U.S. patent application Ser. No. 13/620,363, entitled DISTRIBUTED ROUTING ARCHITECTURE, and filed Sep. 14, 2012, which is a continuation of U.S. patent application Ser. No. 12/641,260, entitled DISTRIBUTED ROUTING ARCHITECTURE, and filed Dec. 17, 2009, now U.S. Pat. No. 8,331,371, the entireties of which are incorporated herein by reference.

BACKGROUND

Generally described, computing devices utilize a communication network, or a series of communication networks, to exchange data. In a common embodiment, data to be exchanged is divided into a series of packets that can be transmitted between a sending computing device and a recipient computing device. In general, each packet can be considered to include two primary components, namely, control information and payload data. The control information corresponds to information utilized by one or more communication networks to deliver the payload data. For example, control information can include source and destination network addresses, error detection codes, and packet sequencing identification, and the like. Typically, control information is found in packet headers and trailers included within the packet and adjacent to the payload data.

In practice, in a packet-switched communication network, packets are transmitted between multiple physical networks, or sub-networks. Generally, the physical networks include a number of hardware devices that receive packets from a source network component and forward the packet to a recipient network component. The packet routing hardware devices are typically referred to as routers. Generally described, routers can operate with two primary functions or planes. The first function corresponds to a control plane, in which the router learns the set of outgoing interfaces that are most appropriate for forwarding received packets to specific destinations. The second function is a forwarding plane, in which the router sends the received packet to an outbound interface.

To execute the control plane functionality, routers can maintain a forwarding information base ("FIB") that identifies, among other packet attribute information, destination information for at least a subset of possible network addresses, such as Internet Protocol ("IP") addresses. In a typical embodiment, the FIB corresponds to a table of values specifying network forwarding information for the router. In one aspect, commercial level routing hardware components can include customized chipsets, memory components, and software that allows a single router to support millions of entries in the FIB. However, such commercial level routing hardware components are typically very expensive and often require extensive customization. In another aspect, commodity-based routing hardware components are made of more generic components and can be less expensive than commercial level routing hardware components by a significant order of magnitude. However, such commodity-based routing hardware components typically only support FIBs on the order of thousands of entries.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing aspects and many of the attendant advantages of this invention will become more readily appreciated

2

as the same become better understood by reference to the following detailed description, when taken in conjunction with the accompanying drawings, wherein:

FIG. 1A is a block diagram illustrative of one embodiment of a distributed routing environment including a router management component and a hierarchical, distributed routing component architecture;

FIG. 1B is a block diagram illustrative of components of a router component utilized in accordance with the distributed routing environment of FIG. 1A;

FIGS. 2A-2C are block diagrams illustrative of the distributed routing environment of FIG. 1A illustrating the routing of a received packet within the hierarchical distributed routing component architecture;

FIG. 3 is a flow diagram illustrative of a distributed router architecture routing routine implemented within a distributed routing environment; and

FIG. 4 is a flow diagram illustrative of a distributed router architecture routing routine implemented within a distributed routing environment.

DETAILED DESCRIPTION

Generally described, the present disclosure corresponds to a distributed routing architecture. Specifically, the present disclosure corresponds to a hierarchical distributed routing architecture including at least three logical levels, or layers, for receiving, processing and forwarding data packets between network components. In one embodiment, the three logical levels can correspond to a core level, a distribution level and a transit level. Illustratively, the core level corresponds to one or more router components that receive an incoming packet from a network component and processes the destination address information associated with the received packet. The core level router component then identifies a distribution level router component based on a subset of the destination address associated with the received packet. The distribution level corresponds to one or more router components that receive a forwarded packet from a core level router component and further processes the destination address information associated with the received packet. The distribution level router component identifies a transit level router component based on at least a subset of the destination address associated with the received packet. Each distribution level router component is associated with, or otherwise corresponds to, a subset of the FIB associated with the distributed routing architecture. Finally, the transit level router components correspond to one or more router components that receive the forwarded packet from a distribution level router component and forward the packet "upstream" to a respective network, or network node. The mapping, or other assignment, of portions of the FIB associated with the distributed routing environment is managed by a router management component.

In one embodiment, each of the router components associated with the core level, distribution level and transit level can correspond more closely to commodity based router components/hardware. In another embodiment, the core level, distribution level and transit level router components correspond to logical router components that do not necessarily have a corresponding hardware router component. For example, one or more logical router components within each level may be implemented in the same hardware router component. Likewise, the logical router components associated with different levels of the distributed routing architecture may be implemented in the same hardware router component. In both embodiments, however, because responsibility for maintaining the FIB associated with the distributed routing

environment is divided among several router components, the processing and memory restraints associated with commodity based router components/hardware can be mitigated. Various implementations, combination, and applications for dividing the FIB associated with the distributed routing environment will be described in accordance with the distributed routing environment. However, one skilled in the relevant art will appreciate that such embodiment and examples are illustrative in nature and should not be construed as limiting.

Turning now to FIG. 1A, a distributed routing environment **100** for implemented a hierarchical distributed routing architecture will be described. The distributed routing environment **100** includes a router management component **102** for controlling the routing information utilized by the distributed routing environment **100**. Specifically, the router managed component **102** can receive all upstream routing information to be used by the distributed routing environment **100** and allocate the assignment of the upstream routing information among the components of the distributed routing environment **100** as will be described. In one embodiment, the router management component **102** can correspond to a computing device in communication with one or more components of the distributed routing environment **100**. Illustrative computing devices can include server computing devices, personal computing devices or other computing devices that include a processor, memory and other components for executing instructions associated with the function of the router management component **102**. In another embodiment, the router management component **102** may be implemented as a software component that is executed on one or more of the router components described below. Illustratively, the router management component **102** maintains and updates the FIB associated with the distributed routing environment **100**. Additionally, the router management component **102** can allocate responsibility for portions of the FIB entries to the various layers of the distributed routing environment **100**, as will be described below. In one embodiment, the router management component **102** can partition the FIB according to the distribution to the various router components of the distributed routing environment **100** and distribute respective portions of the FIB to be maintained in a memory associated with the various router components.

With continued reference to FIG. 1A, the distributed routing environment **100** includes a first communication network **104** that transmits data packets to the distributed routing environment **100**. The first communication network **104** may encompass any suitable combination of networking hardware and protocols necessary to establish packet-based communications to the distributed routing environment **100**. For example, the communication network **104** may include private networks such as local area networks (LANs) or wide area networks (WANs) as well as public or private wireless networks. In such an embodiment, the communication network **104** may include the hardware (e.g., modems, routers, switches, load balancers, proxy servers, etc.) and software (e.g., protocol stacks, accounting software, firewall/security software, etc.) necessary to establish a networking link with the distributed routing environment **100**. Additionally, the communication network **104** may implement one of various communication protocols for transmitting data between computing devices. As will be explained in greater detail below, the communication protocols can include protocols that define packet flow information, such as network address information corresponding to the Internet Protocol version 4 (IPv4) and the Internet Protocol version 6 (IPv6) Internet Layer communication network protocols. One skilled in the relevant art will appreciate, however, that present disclosure

may be applicable with additional or alternative protocols and that the illustrated examples should not be construed as limiting.

In communication with the first communication network **104** is a first level of the distributed routing environment **100**, generally referred to as the core layer or core level. In one embodiment, the core level corresponds to one or more logical router components, generally referred to as core level routers **106A**, **106B**, and **106C**. As previously described, within the distributed routing environment **100**, the core level routers **106A**, **106B**, **106C** receive an incoming packet from a component from the network **104** and process the destination address by identifying a distribution level router component based on a subset of the destination address associated with the received packet. Illustratively, the subset of the destination address can correspond to less than the entire destination IP address, such as the highest most values of the IP address. As previously described, the core level routers **106A**, **106B**, **106C** can correspond to logical router components implemented on one or more hardware components. In one embodiment, each logical router component can correspond with a dedicated physical router component. In another embodiment, each logical router component can correspond to a physical router component shared by at least one other logical router component in the distributed router environment **100**. In an alternative embodiment, at least some portion of the core layer may be implemented by components outside the distributed routing environment **100**. In such an embodiment, such external components would directly address a distribution level router component (described below) of the distributed routing environment **100**.

The distributed routing environment **100** can further include a second level of logical router components, generally referred to as the distribution layer or distribution level. In one embodiment, the distribution level corresponds to one or more router components, generally referred to as distribution level routers **108A**, **108B**, and **108C**. As previously described, within the distributed routing environment **100** the distribution level routers **108A**, **108B** and **108C** receiving an incoming packet from a core routing component **102** and process the destination address by identifying a transit level router component based on at least a subset of the destination address associated with the received packet. Illustratively, the subset of the destination address can correspond to a larger subset of the destination IP address used by the core level routers **106A**, **106B**, **106C**. In this embodiment, the routing performed by the distribution level can correspond to a more refined routing of the received packet relative to the core level routing. As described above with the core level routers **106A**, **106B**, **106C**, the distribution level routers **108A**, **108B**, and **108C** can correspond to logical router components implemented on one or more hardware components. In one embodiment, each logical router component can correspond with a dedicated physical router component. In another embodiment, each logical router component can correspond to a physical router component shared by at least one other logical router component in the distributed router environment **100**.

In communication with the distribution level router components is a third level of router components, generally referred to as the transit layer or transit level. In one embodiment, the transit level corresponds to one or more router components, generally referred to as transit level routers **110A**, **110B**, and **110C**. As previously described, the transit level routers **110A**, **110B**, **110C** receive the forwarded packet from a distribution level router component **108A**, **108B**, **108C** and forward the packet "upstream" to another communication network **112** node. Illustratively, each transit

5

level router **110A**, **110B**, **110C** can be configured to communicate with one or more upstream peers such that all packets destined for an associated peer network component will be transmitted through the assigned transit level router **110A**, **110B**, **110C** (or a redundant router). As described above with the core level routers **106A**, **106B**, **106C** and the distribution level routers **108A**, **108B** and **108C**, the transit level routers **110A**, **110B**, and **110C** can correspond to logical router components implemented on one or more hardware components. In one embodiment, each logical router component can correspond with a dedicated physical router component. In another embodiment, each logical router component can correspond to a physical router component shared by at least one other logical router component in the distributed routing environment **100**.

Similar to communication network **102**, communication network **112** may encompass any suitable combination of networking hardware and protocols necessary to establish packet-based communications to the distributed routing environment **100**. For example, the communication network **112** may include private networks such as local area networks (LANs) or wide area networks (WANs) as well as public or private wireless networks. In such an embodiment, the communication network **112** may include the hardware (e.g., modems, routers, switches, load balancers, proxy servers, etc.) and software (e.g., protocol stacks, accounting software, firewall/security software, etc.) necessary to establish a networking link with the distributed routing environment **100**. As described above with regard to the communication network **104**, the communication network **112** may implement one of various communication protocols for transmitting data between computing devices. One skilled in the relevant art will appreciate, however, that present disclosure may be applicable with additional or alternative protocols and that the illustrated examples should not be construed as limiting.

In an illustrative embodiment, the logical router components (**106**, **108**, **110**) in FIG. 1A may correspond to a computing device having processing resources, memory resources, networking interfaces, and other hardware/software for carrying the described functionality for each of the logical router components. With reference now to FIG. 1B, a block diagram illustrative of components of a router component **150** utilized in accordance with the distributed routing environment **100** of FIG. 1A will be described. The general architecture of the router component **150** depicted in FIG. 1B includes an arrangement of computer hardware and software components that may be used to implement one or more logical router components **106**, **108**, **110**. Those skilled in the art will appreciate that the router component **150** may include many more (or fewer) components than those shown in FIG. 1B. It is not necessary, however, that all of these generally conventional components be shown in order to provide an enabling disclosure.

As illustrated in FIG. 1B, the router component **150** includes a processing unit **152**, at least one network interface **156**, and at least one computer readable medium drive **158**, all of which may communicate with one another by way of a communication bus. The processing unit **152** may thus receive information and instructions from other computing systems or services via a network. The processing unit **152** may also be associated with a first memory component **154** for recalling information utilized in the processing of destination address information, such as at least a portion of a FIB associated with the distributed routing environment **100**. The memory **154** generally includes RAM, ROM and/or other persistent memory. The processing unit **152** may also communicate to and from memory **160**. The network interface

6

156 may provide connectivity to one or more networks or computing systems. The at least one computer readable medium drive **158** can also correspond to RAM, ROM, optical memory, and/or other persistent memory that may persist at least a portion of the FIB associated with the distributed routing environment **100**. In an illustrative embodiment, the access time associated with the memory component **154** may be faster than the access time associated with the computer readable medium driver **158**. Still further, the computer readable medium drive **158** may be implemented in a networked environment in which multiple router components **150** share access to the information persisted on the computer readable medium drive **158**.

The memory **160** contains computer program instructions that the processing unit **152** executes in order to operate the dynamic classifier. The memory **160** generally includes RAM, ROM and/or other persistent memory. The memory **160** may store an operating system **162** that provides computer program instructions for use by the processing unit **152** in the general administration and operation of the router component **150**. The memory **160** may further include computer program instructions and other information for implementing one or more of the logical router components in the distributed routing environment **100**. For example, in one embodiment, the memory **160** includes a router module **164** that implements the functionality associated with any of the routers **106**, **108**, **110**. In the event that multiple logical routers are implemented by the same router component **150**, memory **160** may have each instance of a router module **164**.

In an illustrative embodiment, each router component **150** may be embodied as an individual hardware component for implementing one or more logical routers **106**, **108**, **110**. Alternatively, multiple router components **150** may be grouped and implemented together. For example, each router component **150** may correspond to an application-specific integrated circuit (ASIC) having a processing unit **152**, memory **154** and memory **160** (or other components with similar functionality). The router components **150** may share one or more components, such as the network interface **156** and computer readable medium **158**, via a common communication bus.

With reference now to FIGS. 2A-2C, the processing of receiving packets by the distributed routing environment **100** will be described. With reference first to FIG. 2A, an incoming packet is received from the communication network **104** to a core level router **106**. The core level router **106** that receives the incoming packet may be selected according to a variety of techniques including, but not limited to, load balancing, random selection, round robin, hashing, and other packet distribution techniques. Upon receipt, the core level router **106** processes destination IP address and utilizes a subset of the destination IP address to identify a second level destination router component that will perform a second level of routing. In an illustrative embodiment, the core level router **106** utilizes the most significant bits of the IP address, such as the eight most significant bits of the destination address. The selection of the subset of IP addresses corresponding to a selection of the most significant bits is generally referred to as prefix. For example, selection of the eight most significant bits corresponds to a prefix length of "8." Selection of the sixteen most significant bits corresponds to a prefix length of "16." One skilled in the relevant art will appreciate that the number of bits utilized by the core level router **106** may vary. Additionally, in an alternative embodiment, the core level router **106** may use different methodologies to allocate, or otherwise subdivide, the address space serviced by the distributed routing environment **100**.

Based on the processing of the first subset of the destination address, the core level router **106** forwards the packet to a distribution level router, in this case illustratively **108A**. As previously described, the receiving distribution level router **108A** processes the destination address of the received packet and also utilizes a subset of the destination IP address to identify a third level router component that will forward the packet to a next network destination (outside of the distributed routing environment **100**). Similar to the core level router **106**, the receiving distribution level router can be configured to utilize a selection of the most significant bits of the IP address (e.g., the prefix) to route the packet. In an illustrative embodiment, the prefix used by the distribution level router **108A** is greater than the prefix used by the core level router **106**. Based on the processing by the distribution level router **106A**, the transit level router **110B** receives the forwarded packet and forwards the packet to a designated designation associated with the communication network **112**.

Turning now to FIGS. 2B and 2C, the allocation of IP addresses or subsets of IP addresses within the distributed routing environment **100** will be described. With reference to FIG. 2B, the core level router **106** distributes some portion of the subset of destination IP addresses to distribution level router **108A** (illustrated at **202**). Distribution level router **108A** in turn further distributes the portions of the IP addresses to transit level routers **110A**, **110B**, and **110C** (illustrated at **204**, **206**, and **208**). With reference to FIG. 2C, the core level router **106** distributes a different portion of the subset of destination IP addresses to distribution level router **108B** (illustrated at **210**). Distribution level router **108B** in turn further distributes the portions of the IP addresses to transit level routers **110A** and **110B** (illustrated at **212** and **214**).

In an illustrative embodiment, the router management component **102** (FIG. 1) can allocate responsibility of subsets of IP addresses to the distribution level routers in a variety of manners. In one embodiment, the router management component **102** can allocate responsibility for the entire set of IP addresses in accordance with assignment of IP addresses equally, or substantially equally, among available routers. In this embodiment, each distribution level router **108** becomes responsible for an equal subset of IP addresses or substantially equal if the IP addresses cannot be divided equally. In another embodiment, the router management component **102** can specify specific distribution level router **108** to handle high traffic IP addresses or prefixes. In this example, the entire subset of IP addresses may be custom selected by the router management component **102**. Alternatively, only the subset of IP addresses meeting a traffic threshold may be custom selected with the remaining portions of IP address automatically distributed.

In still a further embodiment, multiple distribution level routers **108** may be selected for a subset of IP addresses. In this embodiment, each core level router **106** can select from multiple distribution level routers **108** based on an equal cost multi-path routing (ECMP) technique in which a specific distribution level router **108** is selected based on a standard load sharing technique. Other factors that can be utilized to select from multiple assigned distribution level router **108** include carrier preference, Internet weather, resource utilization/health reports, an allocated or determine routing cost, service level agreements (SLAs), or other criteria.

In one embodiment, each distribution router **108** can maintain the portion of the FIB that is associated with the subset of IP addresses assigned the respective distribution level router **108**. In another embodiment, each distribution level router **108** can maintain the entire FIB associated with the distrib-

uted routing environment **100** in a memory component, such as computer readable medium **158** (FIG. 1B). Once a subset of IP addresses are assigned to each respective distribution level router **108** (or otherwise updated), the applicable portions of the FIB are loaded in a different memory components, such as memory component **154** (FIG. 1B) utilized by the router (e.g., a routing chip level content addressable memory or a processor level cache memory). The maintenance of the applicable portions of the FIB in a memory component facilitates better router performance by faster memory access times for the applicable portion of the FIB. However, in this embodiment, the allocation of FIBs to each distribution level router **108** can be modified by loading different portions of the stored FIB from a first memory component storing the entire FIB (e.g., the computer readable medium **158**) to the memory component maintaining the portion of the FIB allocated to the distribution level router **108** (e.g., memory component **154**). Accordingly, this embodiment facilitates the dynamic allocation of distribution level routers **108**, the creation of redundant distribution level routers, and additional failover for distribution level routers. Additionally, one or more core level routers **106** can utilize a similar technique in performing the functions associated with the core level of the distributed routing environment **100**.

In still a further embodiment, as a variation to the above embodiment, each distribution level router can be allocated a larger portion of the FIB associated with the distributed routing environment **100** than is capable of being maintained in a first memory component of the router, such as memory component **154** (e.g., a processor level cache memory). If a core level router **106** routes to a distribution level router **108** and the corresponding prefixes of the destination IP address do not correspond to the FIB maintained in the first memory component of the distribution level router, the distribution level router can recall the necessary information from the larger subset of the FIB maintained in a different memory component (e.g., computer readable medium **158** (FIG. 1B)). The FIB maintained in the first memory component (e.g., memory component **152**) may be updated to store the prefix in the primary memory component. Alternatively, the FIB in the first memory component may not be automatically updated based on a single request, but based on increases in traffic for a given prefix.

In yet another embodiment, lower traffic prefixes may be assigned to multiple distribution level routers **108**. In one example, each assigned distribution level router **108** does not maintain the lower traffic routing portion of the assigned FIB in the primary memory component. Rather, routing requests for the lower traffic prefixes can be directed to a specific distribution level router based on selection techniques, such as ECMP, and can be processed by a selected distribution level router **108** based on the larger FIB maintained in a different memory component within the selected distribution level router.

With reference now to FIG. 3, a routine **300** for routing packets and implemented in a distributed routing environment **100** will be described. At block **302**, the distributed routing environment **100** obtains a routing request. As previously described, the routing request is received from a first network **102** (FIG. 1) and includes information identifying a destination IP address. At block **304**, a core level router **106** corresponding to a first level of the distributed routing environment **100** is selected and receives the routing request. In an illustrative embodiment, each core level router **106** can perform the same function and can be selected in accordance with

standard selection techniques, including, but not limited to, random selection, round robin selection, load balancing selection and the like.

At block 306, the selected core level router 106 identifies a distribution level router 108 corresponding to a second level of the distributed routing environment 100. The core level router 106 selects the distribution level router 108 based on processing the destination IP address and utilizing a subset of the destination IP addresses (e.g., the prefix) to determine the appropriate distribution level router 108. Illustratively, in accordance with an embodiment corresponding to the IPv4 communication protocol, the core level router 106 processing can be based on consideration of a prefix of the eight most significant bits. At block 308, the selected distribution level router 108 identifies a transit level router 110 based on processing the destination IP address and utilizing a subset of the destination IP address to determine the appropriate transit level router 110. Illustratively, in accordance with an embodiment corresponding to the IPv4 communication protocol, the distribution level router 108 processing can be based on a larger subset of IP address (e.g., a longer prefix such as 16 or 24 bits, as needed to select the appropriate transit level router 110). One skilled in the relevant art will appreciate, however, the blocks 306 and 308 may be implemented in a manner such the core level router 106 and distribution level router 108 may utilize additional or alternative attributes (including different portions of a destination IP address) of received packets in identifying the next router component to forward the received packet.

At block 310, the selected transit level router 110 transmits the receive packet to the destination recipient associated, or otherwise configured, with the transit level router 110. At block 312, the routine 300 terminates.

With reference now to FIG. 4, another routine 400 for routing packets and implemented in a distributed routing environment 100 will be described. In an illustrative embodiment, routine 400 may be implemented in embodiments in which less than all the FIB associated with a particular distribution router 108 is maintained in a primary memory component. At block 402, a routing request is received at a distribution level router 108. The selection and routing to a distribution level router 108 was previously described above. Although routine 400 will be described with regard to implementation by a distribution level router 108, one skilled in the relevant art will appreciate that at least portions of routine 400 may be implemented by other components of the distributed routing environment 100, such as core level routers 106 or transit level router 110. At decision block 404, a test is conducted to determine whether the subset of the destination IP address associated with the routing request is in the portion of the FIB table maintained in the primary memory of the selected distribution level router 108. If so, at block 406, the distribution level router 108 obtains the transit layer routing information from the FIB maintained in the first memory component (e.g., memory component 152 (FIG. 1B)). At block 408, the distribution level router 108 forwards the packet to the selected transit level router 110.

Alternatively, if at decision block 404 the subset of the destination IP address associated with the routing request is not maintained in the portion of the FIB table maintained in the primary memory of the selected distribution level router 108, at block 410, distribution level router 108 attempts to obtain additional transit routing information from a separate memory component associated with the distribution level router. At block 410, the distribution level router 108 can update the forwarding table information maintained in the primary memory component with the information obtained

from the other memory component. Alternatively, block 410 can be omitted or is otherwise optional. At block 412, the routine terminates.

While illustrative embodiments have been disclosed and discussed, one skilled in the relevant art will appreciate that additional or alternative embodiments may be implemented within the spirit and scope of the present disclosure. Additionally, although many embodiments have been indicated as illustrative, one skilled in the relevant art will appreciate that the illustrative embodiments do not need to be combined or implemented together. As such, some illustrative embodiments do not need to be utilized or implemented in accordance with the scope of variations to the present disclosure.

Conditional language, such as, among others, “can,” “could,” “might,” or “may,” unless specifically stated otherwise, or otherwise understood within the context as used, is generally intended to convey that certain embodiments include, while other embodiments do not include, certain features, elements, or steps. Thus, such conditional language is not generally intended to imply that features, elements or steps are in any way required for one or more embodiments or that one or more embodiments necessarily include logic for deciding, with or without user input or prompting, whether these features, elements or steps are included or are to be performed in any particular embodiment. Moreover, unless specifically stated otherwise, or otherwise understood within the context as used, is generally intended to convey utilization of the conjunction “or” in enumerating a list of elements does not limit the selection of only a single element and can include the combination of two or more elements.

Any process descriptions, elements, or blocks in the flow diagrams described herein and/or depicted in the attached figures should be understood as potentially representing modules, segments, or portions of code which include one or more executable instructions for implementing specific logical functions or steps in the process. Alternate implementations are included within the scope of the embodiments described herein in which elements or functions may be deleted, executed out of order from that shown or discussed, including substantially concurrently or in reverse order, depending on the functionality involved, as would be understood by those skilled in the art. It will further be appreciated that the data and/or components described above may be stored on a computer-readable medium and loaded into memory of the computing device using a drive mechanism associated with a computer-readable medium storing the computer executable components, such as a CD-ROM, DVD-ROM, or network interface. Further, the component and/or data can be included in a single device or distributed in any manner. Accordingly, general purpose computing devices may be configured to implement the processes, algorithms, and methodology of the present disclosure with the processing and/or execution of the various data and/or components described above. Alternatively, some or all of the methods described herein may alternatively be embodied in specialized computer hardware. In addition, the components referred to herein may be implemented in hardware, software, firmware or a combination thereof.

It should be emphasized that many variations and modifications may be made to the above-described embodiments, the elements of which are to be understood as being among other acceptable examples. All such modifications and variations are intended to be included herein within the scope of this disclosure and protected by the following claims.

11

What is claimed is:

1. A system for routing packets, the system comprising:
one or more routers associated with a first level of a hierarchy of physical routing components;
one or more routers associated with a second level of the hierarchy of physical routing components; and
a management component for distribution of forwarding information associated with the first and second levels of the hierarchy, the management component comprising one or more processors configured with specific computer-executable instructions to:
receive the forwarding information for routing data packets addressed to individual addresses within a set of network addresses, wherein the set of network addresses are expressed in a network address format;
select a first subset of the network address format to associate with the first level of the hierarchy, wherein the first subset of the network address format identifies a specific segment of the individual addresses of the set of network addresses;
assign to a router of the first level a first portion of the set of network addresses, wherein the first portion of the set of network addresses is defined at least in part by the first subset of the set of network address format;
distribute, to the determined router of the first level, a first section of the forwarding information, the first section including routing information for the first portion of the set of network addresses;
select a second subset of the network address format to associate with the second level of the hierarchy, wherein the second subset of the network address format identifies a second specific segment of the individual addresses of the set of network addresses;
assign to a router of the second level a second portion of the set of network addresses, wherein the second portion of the set of network addresses is defined at least in part by the second subset of the set of network address format; and
distribute, to the router of the second level, a second section of the forwarding information, the second section including routing information for the second portion of the set of network addresses.
2. The system of claim 1, wherein the network address format corresponds to at least one of an internet protocol version four (IPv4) address format or an internet protocol version six (IPv6) address format.
3. The system of claim 1, wherein the first subset of the network address format corresponds to a first eight or sixteen bits of a 32-bit address format.
4. The system of claim 1, wherein the second subset of the network address format corresponds to a first sixteen or twenty four bits of a 32-bit address format.
5. The system of claim 1, wherein the first portion of the set of network addresses is determined based at least in part on a traffic volume associated with the first portion of the set of network addresses.
6. The system of claim 1, wherein the router of the second level is configured to utilize the second section of the forwarding information to determine a destination component for one or more received packets.
7. The system of claim 6, wherein the destination component is at least one of a component included within a third level of the hierarchy of physical routing components or a component external to the hierarchy of physical routing components.

12

8. A computer-implemented method for managing packet routing within a hierarchy of physical routing components, the computer-implemented method comprising:

- associating a first subset of a network address format with a first level of the hierarchy of physical routing components, wherein the network address format corresponds to a set of network addresses expressible within the network address format, and wherein the first subset of the network address format identifies a specific segment of individual addresses within the set of addresses;
 - assigning to a router of the first level a first portion of the set of network addresses, wherein the first portion of the set of network addresses is defined at least in part by the first subset of the network address format;
 - distributing a first section of a set of forwarding information to the router of the first level, wherein the forwarding information includes information for routing data packets addressed to at least one of the set of network addresses, and wherein the first section includes routing information for the first portion of the set of network addresses;
 - associating a second subset of the network address format to a second level of the hierarchy of physical routing components, wherein the second subset of the network address format identifies a second specific segment of the individual addresses of the set of network addresses;
 - assigning to a router of the second level a second portion of the set of network addresses, wherein the second portion of the set of network addresses is defined at least in part by the second subset of the network address format; and
 - distributing a second section of the forwarding information to the router of the second level, the second section including routing information for the second portion of the set of network addresses.
9. The computer-implemented method of claim 8, wherein the second subset of the network address format is represented by a greater number of bits than the first subset of the network address format.
 10. The computer-implemented method of claim 8, wherein the forwarding information forms a forwarding information base (FIB).
 11. The computer-implemented method of claim 8 further comprising determining at least two additional portions of the set of network addresses to assign to individual routers of the first level, wherein the at least two portions of the set of network addresses include an equal number of network addresses.
 12. The computer-implemented method of claim 8 further comprising assigning the first portion of the set of network addresses to an additional router of the first level, and distributing the first section of the forwarding information to the additional router of the first level.
 13. The computer-implemented method of claim 8, wherein the router of the first level is configured to store the first section of the forwarding information in a first memory component, and wherein the computer-implemented method further comprises distributing an additional section of the forwarding information to the router of the first level for storage in a second memory component.
 14. The computer-implemented method of claim 8, wherein the first portion of the set of network addresses is determined based at least in part on a traffic volume associated with the first portion of the set of network addresses.
 15. The computer-implemented method of claim 14, wherein the first portion of the set of network addresses is determined based at least in part on a determination that the

13

traffic volume associated with the first portion of the set of network addresses satisfies a threshold traffic volume.

16. Non-transitory computer-readable storage media including computer-executable instructions for managing packet routing within a hierarchy of physical routing components, wherein the computer-executable instructions, when executed by a processor, cause the processor to:

select a first subset of a network address format to associate with a first level of the hierarchy of physical routing components, wherein the network address format corresponds to a set of network addresses expressible within the network address format, and wherein the first subset of the network address format identifies a specific segment of individual addresses of the set of addresses;

determine a first portion of the set of network addresses to assign to a router of the first level, wherein the first portion of the set of network addresses is defined at least in part by the first subset of the network address format;

distribute a first section of a set of forwarding information to the router of the first level, wherein the set of forwarding information includes information for routing data packets addressed to one of the set of network addresses, and wherein the first section includes routing information for the first portion of the set of network addresses;

select a second subset of the network address format to associate with a second level of the hierarchy of physical routing components, wherein the second subset of the network address format identifies a second specific segment of the individual addresses of the set of network addresses;

determine a second portion of the set of network addresses to assign to a router of the second level, wherein the

14

second portion of the set of network addresses is defined at least in part by the second subset of the network address format; and

distribute a second section of the forwarding information to the router of the second level, the second section including routing information for the second portion of the set of network addresses.

17. The non-transitory computer-readable storage media of claim 14, wherein the computer-executable instructions further cause the processor to assign the first portion of the set of network addresses to an additional router of the first level, and to distribute the first section of the forwarding information to the additional router of the first level.

18. The non-transitory computer-readable storage media of claim 14, wherein the router of the first level includes a first memory component and a second memory component, the first memory component associated with a faster access time than the second memory component, and wherein the first section of the forwarding information contains more information than may be simultaneously stored in the first memory component.

19. The non-transitory computer-readable storage media of claim 14, wherein the first portion of the set of network addresses is determined based at least in part on a traffic volume associated with the first portion of the set of network addresses.

20. The non-transitory computer-readable storage media of claim 14, wherein at least two routers of the first level are implemented by a single physical router component.

* * * * *